

METAPHOR, METAPHOR SCHEMATA, AND SELECTIVE INFERENCING

Technical Note 204

December 1979

By: Jerry R. Hobbs
Computer Scientist
Artificial Intelligence Center

SRI Projects 7910 and 7500

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 1979		2. REPORT TYPE		3. DATES COVERED 00-12-1979 to 00-12-1979	
4. TITLE AND SUBTITLE Metaphor, Metaphor Schemata, and Selective Inferencing				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International,333 Ravenswood Avenue,Menlo Park,CA,94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 41	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

METAPHOR, METAPHOR SCHEMATA, AND SELECTIVE INFERENCING

Jerry R. Hobbs
SRI International
Menlo Park, California

Abstract

The importance of spatial and other metaphors is demonstrated. An approach to handling metaphor in a computational framework is described, based on the idea of selective inferencing. Three examples of metaphors are examined in detail in this light -- a simple metaphor, a spatial metaphor schema, and a novel metaphor. Finally, there is a discussion, from this perspective, of the analogical processes that underlie metaphor in this approach and what the approach says about several classical questions about metaphor.

1. Metaphor is Pervasive

I. A. Richards, in speaking of metaphor, said, "Literal language is rare outside the central parts of the sciences." (Richards 1936). But it is rare even in the central parts of the sciences. Consider for example the following text from computer science. It comes from an algorithm description in the first volume of Knuth's Art of Computer Programming, Vol. 1, p. 417, and is at but one remove from the domain's most formal mode of expression.

"Given a pointer PO, this algorithm sets the MARK field to 1 in NODE(PO) and in every other node which can be reached from NODE(PO) by a chain of ALINK and BLINK pointers in nodes with ATOM = MARK = 0. The algorithm uses three pointer variables, T, Q, and P, and modifies the links and control bits during its execution in such a way that all ATOM, ALINK, and BLINK fields are restored to their original settings after completion, although they may be changed temporarily."

In this text, the algorithm, or the processor that executes it, is apparently a purposive agent that can perform such actions as receiving

pointers; setting, changing, and restoring fields; reaching nodes; using variables for some purpose; modifying links and bits; and executing and completing its task.

Nodes are apparently locations which can be linked and strung into paths by pointers and visited by the processor-agent.

Nodes also seem to be containers which can contain fields.

Fields are also containers which can contain pointers, among other things. In addition, fields are entities that can be placed at, or set to, locations on the number scale or in the structured collection of nodes.

Pointers, by their very name, suggest objects that can point to a location for the sake of some agent's information.

In fact, there is very little in the paragraph that does not rest on some spatial or agent metaphor. Moreover, these are not simple isolated metaphors; they are examples of large-scale "metaphor schemata" which we use to encode and organize our knowledge about the objects of computer science. They are so deeply ingrained that their metaphorical character generally escapes our notice.*

The pervasiveness of metaphor was noted as early as the eighteenth century by Jeremy Bentham (cf. Ogden 1932). In our century, this observation has been the basis for a rejection of Aristotle's and Quintillian's views that metaphor is mere ornament, and an elevation of metaphor to an "omnipresent principle of language" (Richards 1936) and "the law of its life" (Langer 1942). Richards argued that metaphor involved complex interactions between two domains, which he called the "tenor", that which is being described, and the "vehicle", that which it is being described in terms of. The tenor is seen in a perspective

* I have occasionally had a computer scientist argue that some of the metaphors, e.g. the "variable as container" metaphor, were not metaphors at all but true descriptions of physical reality. To see that this is not the case, note that when we place a value in a variable, its previous value is no longer there; we did not have to remove it. (I once had a beginning FORTRAN student who was puzzled by this very fact. He had not yet learned the limits of the metaphor.)

provided by the vehicle, either bringing to the fore certain aspects of the tenor or allowing the tenor to be viewed in ways that would not have been possible prior to the metaphor.

As we saw in our example, the spatial metaphor especially is pervasive. Jespersen (1922) remarked on this. For Whorf (1939) it was a key element in his view that language determines thought: the spatial metaphors provided by one's language determines how one will normally conceptualize abstract domains. Urban (1939) saw in the use of originally spatial words for more abstract concepts an "upward movement" of language from the physical to the spiritual. The Fundamental Insight that informs all this work is this: Metaphor is pervasive in everyday discourse and is essential in our conceptualizations of abstract domains.

In the last decade, there have been a number of attempts to accommodate this insight in certain traditional frameworks for the study of language. Clark (1973) examines the physical and psychological motivations behind our most common spatial metaphors for time. In Hobbs (1974, 1976) there is an attempt to incorporate the insight into a computational framework; the present paper continues the attempt. In Jackendoff (1976) we find a similar effort in theoretical linguistics. The most thoroughgoing treatment of metaphor in everyday language is found in Lakoff and Johnson (1979); they identify the core metaphors that underlie our thinking about a vast array of domains, and argue that we can understand the domains only by means of these metaphors.

The most detailed proposal in natural language processing for handling metaphor is that of Russell (1976). Her proposal concerns abstract uses of verbs of motion and involves lifting selectional constraints on the arguments of the verb while keeping fixed the topological properties of the motion, such as source, path and goal. Thus, to handle "the ship plowed through the sea," we lift the restriction on "plow" that the medium be earth and keep the property that the motion is in a substantially straight line through some medium.

This is an example of what might be called a "Katzian" approach to metaphor. Its most complete development can be found in Levin (1977), but it is also exemplified in linguistics in the work of Matthews (1971) and Kahn (1975). Metaphor is treated as a species of semantic deviance; selectional constraints are lifted until the expression can sail through the interpreter without difficulty and without effect. But the problem of interpreting "the ship plowed through the sea" is not to avoid rejecting the sentence because the sea is not earth, but to notice the similarity of the wedge-shaped plow and the wedge-shaped bow of a ship and the wake that each leaves, and perhaps more importantly, to take note of the ship's steady, inexorable progress. Any approach to metaphor that does only the first of these is not a way of interpreting metaphors, only of ignoring them. Under the Katzian view, the Fundamental Insight is simply bizarre and inexplicable.*

The work in artificial intelligence that should be most relevant to a study of metaphor is research on analogical reasoning. There are a number of examples. Evans (1968) wrote a program for solving geometric analogy problems. Kling (1971) built a system for proving theorems in ring theory by examining proofs of analogous theorems in group theory (a class of analogies that forms the basis of Galois theory (cf. Artin 1959)). Dershowitz and Manna (1977) and Moll and Ulrich (1979) attempt the automatic synthesis of programs by analogy with known programs. Most of this work has been conducted at too specific a level to be of use in our work on metaphor. Where the specific domain has been abstracted away from, e.g. in Kling (1971) and J. McDermott (1979), the framework has been too general to offer any new insights. An exception to this is Winston (1978), whose work is mentioned in Part 3 below.**

* For further arguments against the Katzian approach to metaphor, see Nunberg (1978).

In this paper I wish to explore how metaphors and metaphor schemata might be treated in a computational setting, in a way that accommodates the Fundamental Insight. In Part 2, three successively more difficult examples of metaphors are considered -- first a simple metaphor, next a metaphor schema that has become a part of the language, and finally a novel metaphor. The aim is to discover some of what is needed to represent and reason about metaphorical usage; the approach employs predicate calculus within a framework for selective inferencing. In Part 3, several issues of classical interest are examined in light of this approach.

** There has also been work on metaphor by psychologists. Most of this (with some notable exceptions discussed in Part 3) involves reaction-time experiments, in which one makes conjectures about the fine details of internal processing on the basis of rather gross input-output behavior, under the dubious assumption that the experimental setup has neutralized all factors other than the one under consideration. It is probably best at present to leave it to psychologists to assess the value of this work. A good review can be found in Ortony (1979b).

2. Representing and Interpreting Metaphors

2.1. Selective Inferencing

One of the principal thrusts of natural language processing research in the last decade has been to develop systems that allow inferences to be drawn selectively (Hobbs 1976, Joshi & Rosenschein 1976, Grosz 1977, Schank & Abelson 1977, A. Robinson 1978). One reason that such systems are needed is that it is difficult, if not impossible, to axiomatize in a consistent manner any domain more complex than set theory. Workers as early as Collins and Quillian (1971) noticed that it is a very powerful device to allow the following inconsistent set of axioms:

- (1) $\text{bird}(x) \rightarrow \text{fly}(x)$
 $\text{ostrich}(x) \rightarrow \text{bird}(x)$
 $\text{ostrich}(x) \rightarrow \sim \text{fly}(x)$

This is a much more economical representation than replacing the first of these axioms with something like

- (2) $\text{bird}(x) \ \& \ \sim \text{ostrich}(x) \ \& \ \sim \text{penguin}(x) \ \& \ \sim \text{kiwi}(x) \ \& \ \sim \text{emu}(x)$
 $\ \& \ \dots \ \& \ \sim \text{injured}(\text{wing}(x)) \ \& \ \sim \text{dead}(x) \ \& \ \sim \text{newborn}(x)$
 $\ \& \ \dots \rightarrow \text{fly}(x)$

The idea is that one can draw an inference as long as it does not result in an inconsistency, and that when an inconsistency does result, some means must be applied to decide among inconsistent inferences. McDermott & Doyle (1978) are developing a nonmonotonic logic in which the various exceptions of (2) are encoded with a special operator M, meaning "it is not inconsistent to assume that". Thus, (2) would be written

$$\text{bird}(x) \ \& \ M \ \text{fly}(x) \rightarrow \text{fly}(x).$$

That is, if x is a bird and it is not inconsistent that x flies, then x flies. For the purposes of this paper, it will be more convenient to

keep the simple notation of (1) and complicate the calculus that manipulates it.

There are further reasons beyond the avoidance of inconsistency to be selective in the inferences one draws -- there are too many true inferences that can be drawn in a specific situation and most of them are irrelevant. Consider the following text:

John couldn't find Mary's house. He drove up one
street and down another.

Among the inferences relevant to understanding this text are the facts that

Houses are visible objects.
Houses are located on streets.
People live in houses.

There are many more facts however that never occur to people in trying to understand this text, and should not "occur" to a natural language processing system. For example:

Houses have roofs.
A House has a living room, a kitchen, several bedrooms, and
one or more bathrooms.
Houses contain furniture.
Houses are made of such materials as wood, brick, stucco.
Termites sometimes attack wooden parts of houses.
Houses have exterior faucets to which hoses can be attached.
Houses tend to rise in value.

All of these things may be true of Mary's house, but the text does not call them to mind.

A great deal of work in natural language processing has been on the problem of using the discourse itself to determine which inferences are relevant. For example, Grosz (1977) examines the clues in task-oriented dialogs that signal a shift to another part of the task and hence another part of the knowledge base. Mann, Moore and Levin (1977) use what is explicit in an utterance to choose a "dialog game", and then assume that what is encoded in that dialog game is relevant. Work on

frame or script recognition (e.g. Schank & Abelson 1977) is in a similar vein. In my own work (e.g. Hobbs 1976) I have investigated the idea that what is relevant are the inferences required to solve various discourse problems, like definite noun phrase resolution. To take a simple example, consider

(3) John sat down in a chair.

It is sometimes true that chairs have arms, and sometimes it is relevant. But there is no reason, given (3) alone, that we would necessarily want to draw the inference that John's chair has arms. However, if the next sentence in the text is

He threw his leg over one of the arms,

then we can be sure that the inference is both true and relevant.

For the rest of the paper we will be assuming a natural language processing system, such as that described in Hobbs (1976), which accepts a text translated by a syntactic front-end into predicate calculus formulae and draws those inferences necessary to solve the discourse problems posed by the text. The inferencing process is selective and driven by a collection of discourse operations which try to do such things as resolve pronoun and definite noun phrase references, find the specific interpretations of general predicates in context ("predicate interpretation"), reconstruct the implicit relation between the nouns in compound nominals, and recognize coherence relations between successive portions of the text. The operations select inferences from a large collection of axioms representing knowledge of the world and the language. Associated with the potential inferences are measures of salience which change as the context changes. These help determine which inferences are drawn by the operations and hence how the text is interpreted.

Three points should be emphasized before we proceed: First of all, we will be dealing at all times at the conceptual level, not at the surface linguistic level. At the conceptual level, we talk about

"predicates", not "words". Metaphors operate primarily at the conceptual level. Although we will generally have, for every word, a predicate of the same name, the predicate should not be thought of as exhausting what is conveyed and suggested by the word. Rather, we should think of the word as corresponding to the possible sets of inferences that might be drawn because the word has been used in a particular context. That is, words do more than merely translate into a single expression in a formal notation; they trigger an inference process that could result in any one of a large set of possible expansions in this notation. We have not stripped words of their mysterious quality, but rather translated it into the mystery of choosing the right set of inferences.

The second point: Black (1979) has emphasized the difference between what he calls a metaphor-theme, a metaphorical formula with no surrounding context, and a metaphor-statement, or simply metaphor, which occurs in some context and must be interpreted in that context. It does not make sense to ask about the interpretation of a metaphor outside of context. This point is often invoked as an argument against a particular formal approach. It must be emphasized that the framework outlined above is specifically designed to formalize the notion of context, and to provide a way of interpreting expressions in context.

Finally, a clarification will be useful. There are at least two questions one might ask about a metaphor: What is its meaning? And what processes are involved in its comprehension? The first is a philosophical question, and I'm not sure it's even well-posed. The second is a psychological question and is the one addressed in this paper. More precisely, we will ask what processes could be involved in the interpretation of metaphors by a cognitive system.

2.2. A Simple Metaphor

With this background, let us return to the problem of metaphor, and consider the simplest case. Suppose it were not a cliché to call someone a hog. How would we go about interpreting the sentence

(4) John is a hog.

Let us suppose our initial logical representation for this is

hog(J).

There are a number of things we might infer from the fact that some entity is a hog, among them

```
hog(x) --> fat(x)
hog(x) --> overconsume(x,y),food(y)
hog(x) --> sloppy(x)
hog(x) --> has-four-legs(x)
```

The problem we are faced with in interpreting (4) is the problem we are always faced with in interpreting a text -- determining which inferences it is appropriate to draw from what we've been told. Depending on the situation, we may want to infer "fat(J)" or "overconsume(J,F)" where "food(F)" or simply "overconsume(J,X)" where X is some other quantity, for example a bench both of us are trying to sit on. The inference that John has four legs is presumably rejected because of strong reasons to believe the contrary. One may or may not infer that John is sloppy, depending on context or other factors.

Our approach then is to say that "John is a hog" conveys at least the information that "John is a hog", but that the various inferences that one could draw from the sentence,

Porky is a hog,

are simply not available to us in the case of John. In particular, most of the inferences that correspond to the various features of the visual image evoked by the word "hog" are not appropriate.

One might object that this approach fails to capture our awareness that (4) does not use "hog" in its literal sense. But what does it mean to know that a word is not used literally? In this framework, it is simply to know that most of a large cluster of canonical inferences cannot be drawn in this instance. In particular, for animal metaphors,

inferences about physical characteristics are suppressed, while inferences about stereotypical character traits are selected. Thus, knowledge that a metaphor is being used is higher level knowledge used to control the inferencing. When a metaphor gets tired and becomes a cliché, as in

John is a chicken,

it could be that the old cluster of canonical inferences -- has feathers, clucks -- is no longer necessarily canonical. The word has acquired a new sense.

The inferences associated with the explicit predication in the metaphor (4) are of three classes. There are those inferences that are definitely intended -- for example, in the right context, the inference "fat(J)" from (4). These ground, or establish a firm basis, for the metaphor; they are what warrant it. Then there are those inferences that are definitely not intended and are inappropriate to draw, the disparities, such as "has-four-legs(J)". Finally, there are inferences that lie in-between, such as "sloppy(J)", which may or may not be intended by the speaker and may or may not occur to the listener. Much of the power of a metaphor derives from this third class of inferences -- the other things that are suggested by the metaphor beyond its firm basis. In fact, even the inappropriate inferences of the second class lend power to the metaphor, since the very denial of something suggests its possibility. The calling up and rejection of the image of a hog in interpreting (4) leaves its trace.

2.3. A Spatial Metaphor Schema

Metaphors that tap into our spatial knowledge are especially powerful since our knowledge of spatial relationships is so extensive, so rich, and so heavily used. As soon as the basis for the spatial metaphor is established, then in our thinking about a new domain we can begin to borrow the extensive machinery we have for reasoning about spatial relationships. For example, once I say that

(5) N is at zero,

and interpret it as

(6) The value of N is equal to zero,

then I have tapped into a large network of other possible uses. I can now say

N goes from 1 to 100

to mean

The value of N successively equals integers from 1 to 100.

I can say

N approaches 100

to mean

The difference between 100 and the value of N becomes smaller.

N can now stay at a number, move from one number to another through several others, be between two numbers, be here, be there. Variables can be scattered along an interval, they can follow one another along the number scale, they can be switched. In short, by means of the simple identification of (5) and (6) we have bought into the whole complex of spatial terminology.

In terms of a system for selective inferencing, what we mean when we say that our spatial terminology is an intricate network is that there are a great many axioms that relate the various spatial predicates. The concept of location -- the predicate "at" -- is at the heart of this network because so many of the axioms refer to it. For example, we might define "go" by means of axioms like

$go(x,y,z) \ \& \ at(w1,x,y) \ \& \ at(w2,x,z) \ \rightarrow \ become(w1,w2)$

that is, if x goes from y to z and w1 is the condition of x being at y

and w2 is the condition of x being at z, then there is a change of state, or a "becoming", from w1 to w2. Similarly, part of the meaning of "switch" can be encoded in the axiom

$$\begin{aligned} &\text{switch}(x,y1,y2) \ \& \ \text{at}(w11,y1,z1) \ \& \ \text{at}(w12,y1,z2) \ \& \\ &\quad \text{at}(w21,y2,z1) \ \& \ \text{at}(w22,y2,z2) \ \rightarrow \\ &\quad \text{cause}(x,\text{become}(\text{and}(w11,w22),\text{and}(w12,w21))). \end{aligned}$$

That is, if x switches y1 and y2 and wij is the condition of yi being at zj, then x causes a becoming from the state in which w11 and w22 hold to a state in which w12 and w21 hold.

We were able to establish the metaphor "a variable as an entity at a location" simply by identifying (5) and (6). In our formalism we can establish the metaphor with similar simplicity by proposing the following axiom:

(7) $\text{variable}(x) \ \& \ \text{value}(w,y,x) \ \rightarrow \ \text{at}(w,x,y)$

That is, if x is a variable and w is the condition of y being its value, then w is also the condition of x being at y.

This simple device of identifying "equality" with "being at" gives us entry into an entire metaphor schema. The schema is represented by a collection of axioms that are intricately woven together by their reference to a small set of common predicates. The schema is tapped for metaphorical purposes by means of axioms like (7), enabling us to transfer to one domain the structure of another, more thoroughly understood domain.

A discourse operation, which in Hobbs (1977b) was called predicate interpretation, uses axioms like (7) to arrive at interpretations of certain metaphorical expressions. The idea behind it is that most utterances make very general or ambiguous sorts of predications and that part of the job of comprehension is to determine the very specific or unambiguous meaning that was intended. Thus, someone might make the general statement

I went to Tokyo,
expecting us to be able to interpret it as

I flew to Tokyo in an airplane,
rather than interpreting the going as swimming, sailing, walking, or any of the myriad other manners of going. In the case of (5), we are expected to determine which of the many ways one thing can be at another is intended in this particular case. That is, rather than determining what we can infer from what is said, we try to determine what the speaker had in mind from which he inferred what he said. In terms of our notation, suppose G is a general proposition and S a specific one and

$$S \rightarrow G$$

is an axiom expressing a fact that a speaker and a listener mutually know. The speaker utters G in the expectation that the listener will interpret it as S. The listener must locate and use the axiom to determine the specific interpretation.

In this manner, axiom (7) provides one possible interpretation of (5), in that it specifies one of the many ways in which one thing can be at another, which the speaker may have meant. When a metaphorical use of "go" or "switch" or any of the other spatial predicates is encountered, axiom (7) combines with the axioms defining the spatial predicate in terms of "at" to give us the correct interpretation.

An alternative to this approach might seem to be to infer intended meaning from what was said. We would use axioms not of the form "S \rightarrow G" but of the form

$$G \ \& \ C_1 \ \& \ \dots \ \& \ C_n \rightarrow M$$

where G is the general proposition that is explicitly conveyed, the C_i's are conditions determinable from context, and M is the intended meaning. For interpreting (5), this would require an axiom like

(8) $\text{at}(w,x,y) \ \& \ \text{variable}(x) \rightarrow \text{value}(w,y,x),$

that is, if w is the condition of x being at y and x is a variable, then w is also the condition of y being the value of x . To interpret (5) we would search through all axioms for axioms that, like (8), have "at" in the antecedent, check whether the other conjuncts in the antecedent were true, and if so, conclude that the axiom's consequent was the intended meaning. This would be equivalent to a "discrimination-net" approach to word-sense disambiguation (e.g. Rieger 1978), in which one travels down a tree-like structure, branching one way or the other according to whether some condition holds, until arriving at a unique specific interpretation at the bottom. The difficulty with this approach is that it supposes we could anticipate at the outset all the ways the meaning of a word could be influenced by context. For metaphors we would have to be able to decide beforehand on all the precise conditions leading to each interpretation. It is highly implausible that we could do this for familiar metaphors, and for novel metaphors the whole approach collapses.

As always, there are a number of inferences involving "at" that we would not want to draw in the case of (5). For example, in the blocks world, if BLOCK1 is at location (2,3,0), then it is impossible for BLOCK2 to be at (2,3,0) at the same time. Yet there is no difficulty whatever in two variables being "at" the same value.* Similarly, if a block is at a location, it is probably being held there by friction and gravity. But with variables there is no need to concern ourselves with what holds them at their values. It is probably the case in general that facts of a "topological" character lend themselves to spatial metaphors, and facts of a "physical" character do not.

One might ask at this point whether our predicate "at" refers to physical location or to some abstract notion of location, or more generally, whether a spatial term subject to metaphorical interpretation

* Even in our casual talk about physical reality, the inference is highly dependent on specific circumstances. We are quite comfortable saying that John and Bill are both at the post office.

has a physical sense and an abstract sense, which are distinct from one another. But we are not doing logical semantics; we are trying to devise a notation for the internal language of a cognitive system and to explicate some of the operations on the notation. For such an endeavor, the primary question is not what an expression in the internal language means or refers to, but rather how it connects up with other expressions in the internal language and with the system's various linguistic and sensorimotor faculties. Of particular relevance is how the spatial predicates connect up with a visual component.

For the sake of concreteness, let us assume the following possible interface between the conceptual and visual components: Assume there are certain predicates which are primitive for the visual component -- "AT", "STRAIGHT" and "BLUE" are good candidates -- such that a proposition containing one of these predicates is automatically asserted when certain things are seen. If along with vision, the visual component implements visualization, then asserting propositions with these primitive predicates would automatically cause the corresponding feature to be visualized. The primitive visual predicates are not the same as the conceptual predicates, like "at", "straight" and "blue", which participate in non-imagistic reasoning about spatial relationships and which are linked with the corresponding lexical items in the linguistic component. However, the two classes of predicates are connected by the following sort of axioms:

- (9) AT(x,y) <--> at(x,y)
 STRAIGHT(x) <--> straight(x)
 BLUE(x) <--> blue(x)

All connections from the visual predicates to the linguistic component and to other expressions in the internal conceptual language are effected through the conceptual predicates by means of these axioms. These axioms, like all axioms, are applied selectively. It is thus possible for the corresponding English words to enter into their many metaphorical uses via the conceptual predicates, while the mechanisms for selective inferencing block visualization by blocking the backward

inferences in (9) (although, just like "has-four-legs", they might contribute through their denial to the metaphor's power.)

Assuming such an interface, let us return to the question: Does "at" refer to physical location? It does and it doesn't. It does, in the sense that all the facts that we know about physical location are expressed in terms of it. It doesn't, in that these inferences are not always drawn and in that the axiom (9) mediates between it and the visualization of physical location.

We are thus saved from what might be called the abstraction proposal, which posits distinct "at" predicates, one which is clearly physical, say "atP", and one abstract, say "atA", which covers both the physical and metaphorical cases. There is simply no need for this. Nothing I have described so far for interpreting metaphors requires any machinery beyond what is required for interpreting nonmetaphorical discourse. While the abstraction proposal may have some plausibility for "at", it is somewhat less plausible to have an abstract predicate "hog" covering both John and real porkers, and it is downright outrageous to propose an abstract predicate "ball" that would also cover bills passed by Congress, as would be required by the example of Section 2.4.*

This may be the least inappropriate place to mention a possible practical application to the domain of computer science. There seem to be two kinds of knowledge we have about the objects of computer science. The first kind is the mathematical or logical definitions of the objects, what might be called their logical structure. On this kind of knowledge are based the programming languages available to us today. However, the way we actually talk about the objects is far richer than the logical structures would seem to warrant, and it ought to be our aim to establish a linkage between the logical structure and the way we talk. This section suggests how this might be done. First basic metaphors for the objects could be outlined, for example, "the processor

* Lakoff and Johnson (1979) raise further objections to the abstraction proposal.

as purposive, sentient agent", "a program as a geography", "a data structure as a geography". The basic metaphors could be linked up with the logical structures through implementations of what may be called metaphor-grounding axioms -- axioms like (7). Then high-level English terminology -- call this rich English -- used in talking about the objects, could be linked up with the basic metaphors via lexical decomposition axioms, and hence linked up with the logical structure of the domain. A schematic of this effort is shown in Figure 1. (We will see this schematic repeated in Part 3 at a more general level.)

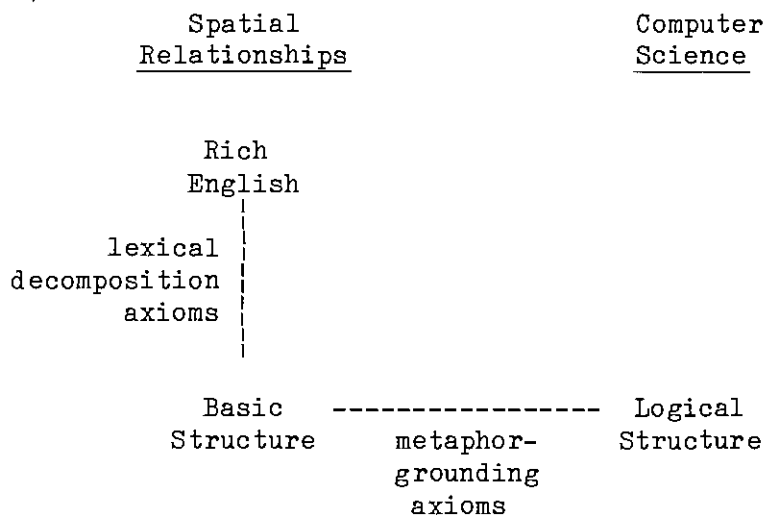


Figure 1.

Our present-day programming languages are very restrictive. It seems desirable to make the rich resources of natural language more available to the programmer as a programming language. Giving programmers access to the domain's most common metaphors would be a significant step in this direction. (This idea is expanded upon in Hobbs 1977a).

2.4. A Novel Metaphor

The final example illustrates how we can represent a metaphor that depends on an elaborate analogy between two complex processes. The metaphor comes from a Newsweek article (July 7, 1975) about Gerald Ford's vetoes of bills Congress has passed. A Democratic congressman complains:

- (10) We insist on serving up these veto pitches that come
over the plate the size of a pumpkin.

It is clear from the rest of the article in which this appears that this means that Congress has been passing bills that the President can easily veto without political damage. There are a number of problems raised by this example, but the only ones we will address are the questions of how to represent and interpret "veto pitches that come over the plate".

The analogy here is between Congress sending a bill to the President to sign or veto and a pitcher throwing a baseball past a batter to miss or hit. Let us try to encode each of the processes first and then to establish the links between them. The facts about a bill are as follows: The participants are Congress, the bill, and the President. Congress sends a bill to the President, who then either signs it or vetoes it. We will assume there is an entity C, Congress. To encode the fact that C is Congress, we could write

Congress(C).

But it will prove more useful to assume there is a condition, call it CC, which is the condition of C being Congress. We will represent this

Congress(CC,C).

Similarly, there are entities B, CB, P, and CP, with the properties

bill(CB,B),

i.e. CB is the condition of B being a bill, and

President(CP,P),

i.e. CP is the condition of P being the President. There are three relevant actions, call them SD, SG, and VT, with the following properties:

send(SD,C,B,P),

i.e. SD is the action by Congress C of sending the bill B to the President P;

sign(SG,P,B),

i.e. SG is the action by the President P of signing the bill B; and

veto(VT,P,B),

i.e. VT is the action by the President P of vetoing the bill B. There is the condition -- call it OSV -- in which either the signing SG takes place or the vetoing VT takes place:

or(OSV,SG,VT).

Finally, there is the situation or condition, TH, of the sending SD happening followed by the alternative actions OSV:

then(TH,SD,OSV).

The corresponding facts about baseball are as follows: * There are a pitcher x, a ball y, and a batter z, and there are the conditions cx, cy, and cz, of x, y, and z being what they are:

pitcher(cx,x)
ball(cy,y)
batter(cz,z).

* Where individual constants, C, CC, B, ..., were used in the Congressional bill schema, universally quantified variables, x, cx, y, ..., are used here. This is because the baseball schema is general knowledge that will be applied to the specific situation involving Congress and the President. It is a collection of axioms that get instantiated in the course of interpreting the metaphor.

The actions are the pitching p by the pitcher x of the ball y to the batter z,

pitch(p,x,y,z);

the missing m of the ball y by the batter z,

miss(m,z,y);

and the hitting h of y by z,

hit(h,y,z).

Let omh represent the condition of one or the other of m and h occurring,

or(omh,m,h),

and th the situation of the pitching p followed by either m or h,

then(th,p,omh).

The linkage established by the metaphor is among other things, between the bill and the ball. But it is not enough to say that B, in addition to being the bill, is also in some sense a ball, just as B has other properties, say, being concerned with federal housing loans, being printed on paper, and containing seventeen subsections. The metaphor is stronger. What the metaphor tells us is that the condition of B being the bill is indeed the condition of B being a ball. Similar links are established among the other participants, actions, and situations. That is, the baseball schema is instantiated with the entities of the Congressional bill schema, leading to the following set of propositions:

(11)	Congress(CC,C)	pitcher(CC,C)
	bill(CB,B)	ball(CB,B)
	President(CP,P)	batter(CP,P)
	send(SD,C,B,P)	pitch(SD,C,B,P)
	sign(SG,P,B)	miss(SG,P,B)
	veto(VT,P,B)	hit(VT,P,B)
	or(OSV,SG,VT)	
	then(TH,SD,OSV)	

The two schemata and their links are shown more graphically in Figure 2.

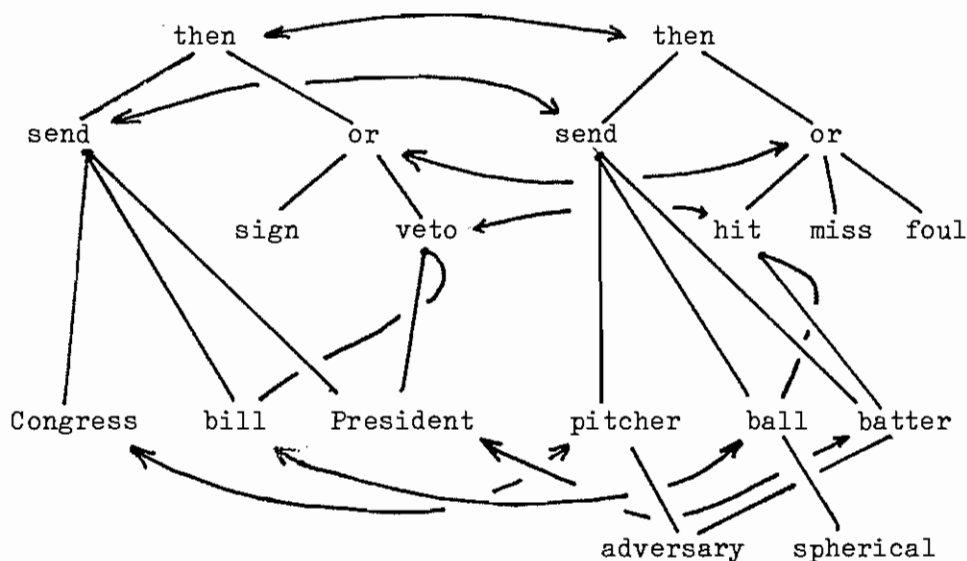


Figure 2.

Although all of this has been described in terms of schemata, a schema in this framework is simply a collection of possibly very complex axioms that are interrelated by the co-occurrence of some of the same predicates, perhaps together with some meta-knowledge for controlling the use of the axioms in inferencing. The linkage between the two schemata does not require some special "schema-mapping" operation, but only the assumption of identity between the corresponding conditions, just as in the second example we identified "equality" with "being at". As in the first two examples, it is because of the mechanisms of selective inferencing that this will do. Thus, to represent the metaphor, we do not have to extend our formalism beyond what was required for the first two examples, nor indeed beyond what is required for nonmetaphorical discourse.

However, a shortcoming of this representation, as it stands, is that there is no explicit separation of the two parts of the metaphor. Thus, C is both Congress and a pitcher and P is both the President and a batter. But there is no explicit indication that the properties

"Congress" and "President" belong to one side of the metaphor and "pitcher" and "batter" to the other. We could remedy this by being more careful about the difference between a condition and a description of the condition. For then we could say that the condition CC of C being Congress is identical to the condition of C being a pitcher, while the descriptions involving "Congress" and "pitcher" are distinct. We would then make assertions about the descriptions that they belong to one domain or the other. But the details of this hastily sketched idea cannot be worked out here.

No natural language processing system existing today could derive (11) from (10). Nevertheless, we can make a reasonable guess as to the basic outline of a solution: The congressman said, "We insist on serving up these veto pitches" For someone to serve up a pitch is for him to pitch. This leads to the identification of Congress with the pitcher. To interpret the compound nominal "veto pitch", we must find the most salient, plausible relation between a veto and a pitch. From our knowledge about vetoes, we know that Congress must first send the bill to the President. From our knowledge about pitching, we know that for the Congress-pitcher to pitch, it must send a "ball" to a "batter". We have a match on the predicate "send" and on the agents of the sendings, Congress. We can complete this match by assuming the bill is the ball and the President is the batter.*

We have almost a complete match between the two situations. The analogy will be completed when we determine which of the various possible actions that a batter can perform corresponds to the President's veto. But this is just what we need to complete the relation between "veto" and "pitch" in the compound nominal. By some means well beyond the scope of this paper to discuss, "pitches that come over the plate the size of a pumpkin" must be interpreted to mean that the ball is easy for the batter to hit. If we assume maximum redundancy -- that a veto pitch and a pitch that comes over the plate the size of a

* Such assumptions are common in interpreting discourse. In fact, they constitute one of the principal mechanisms for resolving pronouns and implicit arguments. (See Hobbs 1979).

pumpkin are roughly the same thing -- then we assume that the pitch is a bill-ball that the Congress-pitcher sends to the President-batter which he then finds easy to veto-hit. The analogy is complete.

As with all metaphorical expressions, as indeed with any expression, there will be a number of inferences we will not want to draw in this case -- for example, that B is spherical and has stitching. But this metaphor invokes other inferences that we do accept, inferences that would not necessarily follow from the facts about the American government. It suggests, for example, that Congress and the President are adversaries in the same way that a pitcher and a batter are, and that from the President's perspective it is good for him to veto a bill Congress has passed and bad for him to sign it. What we know about the adversary relationship in baseball is vivid and unambiguous, and herein lies the power of the metaphor.

This example involves the identification of two highly structured portions of our knowledge base. It raises a question of whether our approach can handle metaphors in which one domain has much less structure, especially metaphors which impart structure to a domain that it would not otherwise have. Lakoff and Johnson (1979) demonstrate this effect by inventing a "love as a collaborative work of art" metaphor and showing some of the things that can be concluded about love as a result. I see no fresh difficulties that this would cause for our approach. Corresponding to the numerous basic links between the existing Congressional bill and baseball schemata, there would be only a few links between our knowledge of love and of collaborative works of art. If this new metaphor is productive, then corresponding to the suggestion from baseball of an adversary relationship in government, there will be numerous suggestions from the nature of collaborative works of art about the nature of love. Therefore, the effect of the new metaphor may be quite different from the effect of the ones we have examined, but the mechanisms involved in interpreting it are the same.

3. Some Questions about Metaphor

In all three examples, we have seen the same broad processes at work. They can be summarized as follows: There are two domains, which we may call the new domain, or the domain which we are seeking to understand or explicate, and the old domain, or the domain in terms of which we are trying to understand the new domain and which provides the metaphor. Richards (1936) refers to these as the tenor and the vehicle, respectively. In our examples the new domains are John's nature, computer science, and the workings of the American government. The old domains are a hog's nature, spatial relationships, and baseball. For each old domain, we can distinguish between what may be called the basic concepts and relationships and complex concepts and relationships. For spatial relationships, "at" is a basic concept; "go", "approach", and "switch" are complex concepts. For baseball, "pitcher" and "batter" are basic, their adversary relationship is complex. What is basic and what is complex in a particular domain are not necessarily fixed beforehand, but may be determined in part by the metaphor itself.

Each of the examples can be viewed as setting up a link between the basic concepts of a new domain and an old domain, in order that complex concepts or relationships will carry over from the old to the new. The following diagram illustrates this:

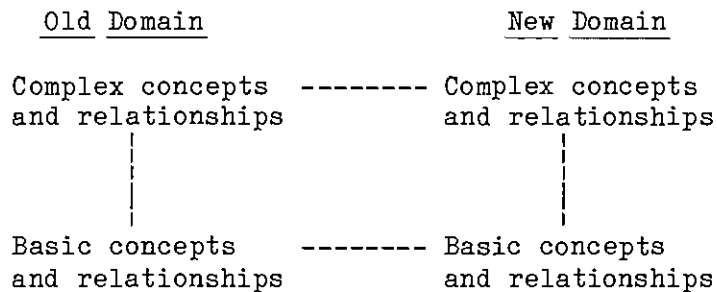


Figure 3.

This diagram is familiar from Galois theory, algebraic topology, and category theory (e.g. Artin 1959, Spanier 1966, MacLane 1971). One can prove theorems in one domain -- for example, the category of fields -- by constructing a "functor" to map its objects and relations into the objects and relations of another domain -- for example, the category of groups -- proving the theorem in the second domain, and using the inverse functor to map it back into the original domain.

The diagram illustrates a general paradigm for analogical reasoning. To reason in a new domain about which we may know little, we map it into an old domain, do the reasoning in the old domain, and map the results back into the new domain. To deal with metaphors in this paradigm we have had to specify the nature of the links between the four corners of the diagram. It has been an argument of this paper that, for the horizontal links, identity will do. For a short time, at least, John really is a hog; "equality" really is "being at"; the bill really is a ball.

What permits this move is that the mechanisms of selective inferencing will take care that the wrong things are not concluded as a result. This brings us to the vertical links in the diagram. They are comprised of the collections of axioms encoding the relationships between basic and complex concepts, such as "lexical decomposition" axioms. But how the links are used now becomes the central problem, for we have translated the question of how to interpret a metaphor into the question of how to choose which inferences are the appropriate ones to draw when a metaphorical expression is encountered in a text. In category theory, once the functor maps the new domain into the old domain, then everything we can conclude in the old domain must carry over to the new. However, in most kinds of analogical reasoning and in interpreting metaphors, only a subset of what can be concluded in the old domain will carry over to the new. The major problem for us, then, is how to determine precisely what subset of facts from the old domain does carry over to the new. Moreover, investigating this question is probably the most useful way of investigating the mechanisms behind metaphor.

A threefold classification of possible inferences is suggested:

1. Groundings, or the inferences that must be drawn if one is to make sense of the metaphor. These are what warrant the metaphor. In our first example, a grounding may be the inference that John overconsumes; in the third example, that the bill-ball is sent to the President-batter.

Black (1962) suggests a classification of theories of metaphor that includes "substitution theories", in which a metaphor is analyzed by replacing the explicit predication with those literal propositions it is intended to convey. * In our terms, it is the grounding inferences that such theorists want to substitute for the metaphor.

2. Disparities, or the inferences that should not be drawn, whether because they are contradictory or irrelevant. In our examples, a disparity between John and a hog is that a hog is a farm animal with four legs and a snout, between the bill and a ball that a ball is spherical.

Richards points out that the disparities frequently play an important role: a significant effect of a metaphor is the recognition that some of the criterial inferences that could be drawn from the explicit predication are not appropriate is an important effect of the metaphor. The fact that John, though a hog, is not a farm animal, but a person, carries the implication that he should resemble a farm animal even less. Ong (1955) points out that a metaphor is effective only as long as it calls these disparities to mind. "John is a hog" strikes us in a way that "the foot of a mountain" does not.

In our approach, certain disparities are considered and actively denied, rejected when inconsistency is discovered. This active process may be compared with a cartoon in which John gradually acquires a snout and hooves while stuffing himself, then returns to his normal appearance. This has the flavor of a "reverse substitution" theory of metaphor, in which the inappropriate properties inferrable from the explicit predication, for a moment, replace the metaphor.

* Beardsley refers to this as the "literalist" theory (1958) and the "comparison" theory (1967).

3. Suggestions, a weak term for one of metaphor's greatest powers, its suggestiveness. These are the inferences that may or may not be drawn. Neither are they required to interpret the metaphor, nor are they obviously inappropriate. In our first example, a suggestion is that John is sloppy, in the third it is suggested that the President and Congress are adversaries.

There are positive and negative aspects to this suggestiveness. On the positive side, it is this more than anything else that makes metaphor such a powerful conceptual tool. By using an apt metaphor to map a new, uncertainly understood domain into an old, well-understood domain, such as spatial relationships, which has a rich collection of axioms connecting the basic and complex levels, we gain a more certain grasp on the new domain conceptually, and linguistically we provide it with a richer vocabulary. We are able to draw conclusions we would not initially have expected.*

On the other hand, Lakoff and Johnson (1979) point out the dangers of mistaking the metaphor for a true description, and thus drawing too many suggested inferences without adequately examining their appropriateness. One is blinded to the limits of the metaphor, and also to alternative metaphors. Reddy (1979) discusses a specific case, the language-as-conduit metaphor and its influence on the study of communication; the theme is also developed at length by Turbayne (1962) with respect to metaphors of science.

We now have several alternative statements of the central question: How are metaphors interpreted? How do we select the appropriate inferences to draw from the use of a metaphorical expression? For each possible inference that could be drawn, how do we determine which of the three classes it falls into on a given occasion? We have only the barest hints of an answer. We should not expect a general, isolated

* Some preliminary knowledge acquisition dialog experiments we have conducted at SRI indicate that to learn about a new domain, one must acquire not only the logical structure of its objects. One must also learn its basic spatial metaphors. So it seems that we try to exploit the suggestive power of metaphor from the outset.

solution. It is more likely that a number of processes will contribute, that the solution will come from a number of directions. There have been several proposals.

Ortony (1979a) has suggested a breakdown of the knowledge in the old and new domains into classificatory facts, other high-salience facts, and low-salience facts. Classificatory facts are not transferred from the vehicle to the tenor. Thus, from "John is a hog" we do not infer that John is a farm animal. What get transferred from the vehicle to the tenor are other high-salience facts whose correlates in the tenor are of low salience. It is a high-salience fact that hogs overconsume, a low-salience fact that John overconsumes. The effect of the metaphor is to bring to the fore this low-salience fact about John.

Winston (1978) presents an algorithm in which properties are transferred from the vehicle to the tenor if they are extremes on some scale, are known to be important, or serve to distinguish the vehicle from other members of its class. Thus, properties of hogs that were not shared by other farm animals would be transferred.

Another possibility was hinted at above. There may be certain classes of inferences that typically get transferred -- stereotypical character traits rather than physical properties for animal metaphors, topological properties rather than physical properties for spatial metaphors, and so on. It is not surprising that this should be the case, since the function of these metaphors is usually to make sense of some abstract domain.

However, none of these approaches take into account the text in which the metaphor is embedded. The principal thesis of this paper has been that the discourse operations a natural language processor must possess anyway -- operations like predicate interpretation, compound nominal interpretation, and coherence relation recognition -- will often serve to pick out the relevant inferences in cases of metaphor. Our second example was interpreted by means of the predicate interpretation operation. The analysis of the third example was driven by the need to interpret the compound nominal "veto pitch". Our first example would

generally be interpreted by trying to relate the utterance to some currently salient portion of our knowledge. I gave insufficient context to allow precise interpretation. But we can make up contexts in which discourse operations become decisive. For example, in the text

Mary eats like a bird, but John is a hog,

coherence considerations (cf. Hobbs 1978) force the interpretation. In order to see "but" as appropriate, we must draw the inference that John overconsumes food.*

While no general method of interpretation has been found, we have made progress just by translating the problem into the problem of selecting the right inferences, for this is the problem of natural language processing. Research on metaphor can now proceed in a way that meshes well with research on other aspects of language understanding. Moreover, we can expect that light will be shed on the problems of metaphor by results of a more general nature in discourse interpretation, in which the metaphorical-nonmetaphorical distinction may not be especially significant.

We are now in a position to address several variants of the problem most authors take up first -- what is metaphor? A computational version of this question might seem to be, how are metaphors recognized? But we will consider a prior question instead: Are metaphors recognized? The standard account is that metaphorical interpretation is triggered by a failure to make sense out of the literal meaning. That is, a separate initial step is postulated in which something is found to be wrong. But this fails to take into account the work that needs to be done to find that everything is right. Even Searle (1979), who discusses at length the difficulties of interpreting literal utterances, separates these processes from the process of interpreting the utterance once the deviance is found, overlooking their likely identity.

* It is important to note that the same problem faces us in these examples as faces us in all of discourse interpretation -- how to navigate in a flexible, efficient manner through the sea of possible inferences.

An excellent example of the difficulties in interpreting literal expressions is provided by what Black (1962) calls the "comparison" view of metaphor. A metaphor is seen as an elliptical form of a simile. Thus, the metaphorical "John is a hog" translates into the literal "John is like a hog" or "John is like the stereotypical hog in certain respects." But the predicate "like" is an very good example of a literal expression whose interpretation is quite problematic. Part of the literal meaning of "A is like B" is that A shares certain properties with B. Thus, in understanding "His house is like my house", we need to determine in which respects the two are alike. Similarly, in interpreting "John is like a hog," we must discover in just what respects John is like a hog. But this means that the problem of interpreting the literal "like" is isomorphic to the problem of interpreting the original metaphor.*

We have seen in this paper that frequently the discourse operations result in a metaphor being interpreted, and that the operations themselves do not depend on the metaphor-nonmetaphor distinction. They are just the ordinary processes of deciding which inferences to draw and which to refrain from drawing. It may be, of course, that once the groundings of the metaphor are discovered, knowledge that it is a metaphor plays a role in directing further inferencing. But metaphor recognition is by no means a computationally necessary part of metaphor interpretation. This also accords with the intuition that understanding metaphorical expressions is not different in kind from understanding nonmetaphorical language.

Miller (1979) revives the comparison view, but to get around Black's charge of vacuity (1962), he analyzes the notion of similarity (cf. also Tversky 1977). The result is something very close to the view presented in this paper. For him, the recognition step is unnecessary, and he gives an excellent account of why. Nunberg (1978) also argues for the identity of the interpretation processes. Rumelhart

* Except of course identity is not assumed between the tenor and the vehicle. This is the standard observation about the difference between metaphor and simile.

(1979) shows that literal interpretation is sometimes problematic, as a way of arguing for the identity of these processes.

Another variant of the question, what is metaphor, is whether to count so-called "tired" or "dead" metaphors as metaphors at all, or to reserve the term for novel examples. Extremes have been argued. Isenberg (1963) urges that "metaphor" be reserved for examples that are not just novel, but have artistic intent. Black (1979) wants to exclude the example "that no longer has pregnant metaphorical use." On the other hand, Richards (1936) and Whorf (1939) see metaphor everywhere -- the Fundamental Insight of Part 1. On the far left, Lakoff and Johnson (1979) even view nominalizations of verbs as examples of an "event-as-object" metaphor.

To get a handle on this issue, let us trace through four stages in the life story of a metaphor, referring again to Figure 3. Think of a novel metaphor as a complex term from the old domain used in a context that requires a concept from the new domain. To interpret it we must decompose the complex term into basic concepts in the old domain, and either use available links between new and old basic concepts or surmise such links for the first time. This enables us to project the complex concept from the old to the new domain. For novel metaphors, we might expect this to require quite a bit of computing, and involve following a number of false leads.

The second stage is when the metaphor has become "familiar". The same path is followed in interpreting it, but now the salience of the required inferences is such that the computation is direct and fast. The path that had to be reconnoitered with some care when the metaphor was novel is now worn into a broad avenue that is difficult not to follow.

In the third stage, the metaphor becomes tired. A direct link is established between the basic and complex levels in the new domain. That is, the expression acquires a new sense, it becomes technical terminology in the new domain. Nevertheless, at this stage, the metaphor can be reactivated (cf. Brooks 1965, Black 1979). We can be

forced to compute anew the path whose computation is no longer ordinarily necessary. For instance, in the same Newsweek article that contains the example of Section 2.4, we find the tired metaphor, "see which way the wind was blowing." This is an idiom, an entry in our lexicon, and normally we need not analyze the comparison it makes between power relationships and meteorological phenomena. But the next sentence is "The wind could change," and this places the comparison squarely before us.

Finally the metaphor dies. Because of changes in the language user's knowledge base or because of the way he learned the expression, he can not recover the path that makes sense of the metaphor. It exists only as an expression in the new domain. Yet at this stage we can still ask, as linguists, what processes "motivate" this expression in this domain (cf. Fillmore 1979), why the expression makes sense, even though as psychologists we do not believe the person uses or could use the processes. An example will help: Suppose someone learns the expression

set a variable to a value,

purely as technical terminology, without ever learning the underlying spatial metaphor of setting a dail to a location. A text that would reactivate the metaphor if it were merely tired only baffles him. The metaphorical nature of the expression cannot be said to play a role in his interpretation of it. Nevertheless, its technical sense is not arbitrary. The technical use of "set to" was originally motivated by the metaphor. The processes used to interpret it when it was novel can be said to motivate it now.

In summary, the four stages can be described thus. In stage 1, the interpretation is computed. In stage 2, it is computed easily. In stage 3, it is computable, though no longer computed; at this stage, reactivation of the metaphor causes it to be computed again. In stage 4, it is neither computed nor computable, but there is nevertheless a "historical" motivation.

Which stages are entitled to be called metaphor? Where should the line be drawn? The above account provides reasons enough for drawing the line anywhere. But in terms of the processes involved, there is simply no point in drawing a line, for they are the same at every stage. What differs is how and when they are used. The reason not to exclude the more decrepit metaphors from our investigation is that they require the same processes to be explicated as do livelier metaphors. But here the processes appear as the processes that motivated the expression, not the processes used to interpret it.

Finally, we can ask what makes a metaphor good. This overlaps with the larger question of what makes any text aesthetically pleasing, and I take this to be a mystery not accessible to present-day cognitive science. Even here, however, a detailed account of the cognitive processes involved in metaphor interpretation may shed some light.

Acknowledgments

I have profited from discussions with Armar Archbold, Gary Hendrix, George Lakoff, and Bob Moore, and also with the participants of the Artificial Intelligence Workshop at the Electrotechnical Laboratory, Tokyo, Japan, in August 1979. None of these people are responsible for any errors in this paper. This work was supported by the National Science Foundation under Grant No. MCS-78-07121 and by the Defense Advanced Research Projects Agency under Contract No. N00039-79-C-0118 with the Naval Electronic Systems Command.

REFERENCES

- Artin, E. 1959. Galois theory. Notre Dame Mathematical Lectures, No. 2. Notre Dame, Indiana.
- Beardsley, M. 1958. Aesthetics: Problems in the philosophy of criticism. New York: Harcourt Brace and Co.
- Beardsley, M. 1967. Metaphor. In P. Edwards (Ed.), Encyclopedia of philosophy, Vol. 5, 284-289.
- Black, M. 1962. Metaphor. In M. Black, Models and metaphors: Studies in language and philosophy, 25-47. Ithaca, New York: Cornell University Press.
- Black, M. 1979. More about metaphor. In A. Ortony (Ed), Metaphor and Thought, 19-43. Cambridge, England: Cambridge University Press.
- Brooks, C. 1965. Metaphor, paradox, and stereotype. British Journal of Aesthetics, Vol. 5, 315-328.
- Clark, H. 1973. Space, time, semantics, and the child. In T. Moore (Ed.), Cognitive development and the acquisition of language. New York: Academic Press.
- Collins, A. and M.R. Quillian. 1971. How to make a language user. In Tulving and Donaldson, Eds. Organization of memory, 309-351.
- Dershowitz, N. and Z. Manna 1977. The evolution of programs: A system for automatic program modification. IEEE Transactions on Software Engineering, Vol. SE-3, No. 4.
- Evans, T. 1968. A program for the solution of a class of geometric-analogy intelligence-test questions. In M. Minsky (Ed.), Semantic Information Processing, 271-353. Cambridge, Mass.: The MIT Press.
- Fillmore, C. 1979. Innocence: A second idealization for linguistics. Proceedings, Fifth Annual Meeting, Berkeley Linguistics Society, 63-76. Berkeley, California.
- Grosz, B. 1977. The representation and use of focus in dialogue understanding. Stanford Research Institute Technical Note 151, Stanford Research Institute, Menlo Park, CA. July 1977.
- Hobbs, J. 1974. A model for natural language semantics, Part I: The model. Research Report 36, Department of Computer Science, Yale University. October 1974.

- Hobbs, J. 1976. A computational approach to discourse analysis. Research Report 76-2, Department of Computer Sciences, City College, City University of New York. December 1976.
- Hobbs, J. 1977a. What the nature of natural language tells us about how to make natural-language-like programming languages more natural. Proceedings, Symposium on Artificial Intelligence and Programming Languages, 85-93. Rochester, N.Y. August 1977.
- Hobbs, J. 1977b. Coherence and interpretation in English texts. Proceedings, International Joint Conference on Artificial Intelligence, 110-116. Cambridge, Mass. August 1977.
- Hobbs, J. 1978. Why is discourse coherent? SRI Technical Note 176, SRI International, Menlo Park, California. November 1978. To appear in F. Neubauer (Ed.) Coherence in natural language texts.
- Hobbs, J. 1979. Coherence and coreference. Cognitive Science, Vol. 3, No. 1, 67-90.
- Isenberg, A. 1963. On defining metaphor. The Journal of Philosophy, Vol. 60, 609-622.
- Jackendoff, R. 1976. Toward an explanatory semantic representation. Linguistic Inquiry, Vol. 7, No. 1, 89-150. Winter 1976.
- Jespersen, O. 1922. Language: Its nature, development and origin. London: George Allen & Unwin Ltd.
- Joshi, A., & S. Rosenschein. 1976. Some problems of inferencing: Relation of inferencing to decomposition of predicates. Statistical Methods in Linguistics, Vol. 1, No. 1.
- Kahn, E. 1975. Frame semantics for motion verbs with application to metaphor. In C. Cogen et al (Eds.), Proceedings of the First Annual Meeting of the Berkeley Linguistic Society, 246-256.
- Kling, R. 1971. A paradigm for reasoning by analogy. Artificial Intelligence, Vol. 2, 147-178.
- Knuth, D. 1973. The art of computer programming, Vol.1. Reading, Mass.: Addison-Wesley.
- Lakoff, G. and M. Johnson 1979. Metaphors we live by.
- Langer, S. 1942. Philosophy in a new key: A study in the symbolism of reason, rite, and art. Cambridge, Mass.: Harvard University Press.
- Levin, S. 1977. The semantics of metaphor. Baltimore: The Johns Hopkins University Press.

- MacLane, S. 1971. Categories for the working mathematician. New York: Springer-Verlag.
- Mann, W., J. Moore and J. Levin 1977. A comprehension model for human dialogue. Proceedings, International Joint Conference on Artificial Intelligence, 77-87, Cambridge, Mass. August 1977.
- Matthews, R. 1971. Concerning a 'linguistic theory' of metaphor. Foundations of Language, Vol. 7, 413-425.
- McDermott, D. and J. Doyle 1978. Non-monotonic logic I, A.I. Memo 486, Artificial Intelligence Laboratory, Massachusetts Institute of Technology. August 1978.
- McDermott, J. 1979. Learning to use analogies. Proceedings, International Joint Conference on Artificial Intelligence, 568-576. Tokyo, Japan. August 1979.
- Miller, G. 1979. Images and models, similes and metaphors. In A. Ortony (Ed), Metaphor and Thought, 202-250. Cambridge, England: Cambridge University Press.
- Moll, R. and J. Ulrich 1979. The synthesis of programs by analogy. To appear in Artificial Intelligence.
- Nunberg, G. 1978. The pragmatics of reference. Ph.D. thesis. City University of New York, New York, NY.
- Ogden, C. 1932. Bentham's theory of fictions. London: Routledge & Kegan Paul Ltd.
- Ong, W. 1955. Metaphor and the twinned vision. Sewanee Review, Vol. 63, 193-201. Spring 1955.
- Ortony, A. 1979a. The role of similarity in similes and metaphors. In A. Ortony (Ed), Metaphor and Thought. Cambridge, England: Cambridge University Press.
- Ortony, A. 1979b. Metaphor. To appear in R. Spiro et al (Eds.), Theoretical issues in reading comprehension. Hillsdale, New Jersey: Erlbaum.
- Reddy, M. 1979. The conduit metaphor -- A case of frame conflict in our language about language. In A. Ortony (Ed), Metaphor and Thought, 284-324. Cambridge, England: Cambridge University Press.
- Richards, I. A. 1936. The Philosophy of Rhetoric. Oxford: Oxford University Press.
- Rieger, C. 1978. Viewing parsing as word sense discrimination. In W. Dingwall (Ed.), A Survey of Linguistic Science. Greylock Pub.

- Robinson, A. 1978. Investigating the process of natural-language communication: A status report. SRI Technical Note 165. SRI International, Menlo Park, California. July 1978.
- Rumelhart, D. 1979. Some problems with the notion of literal meanings. In A. Ortony (Ed), Metaphor and Thought, 78-90. Cambridge, England: Cambridge University Press.
- Russell, S.W. 1976. Computer understanding of metaphorically used verbs. American Journal of Computational Linguistics, microfiche 44.
- Schank, R. and R. Abelson 1977. Scripts, plans, goals, and understanding. Hillsdale N.J.: Laurence Erlbaum Associates.
- Searle, J. 1979. Metaphor. In A. Ortony (Ed), Metaphor and Thought, 92-123. Cambridge, England: Cambridge University Press.
- Spanier, E. 1966. Algebraic topology. New York: McGraw Hill.
- Turbayne, C. 1962. The myth of metaphor. Columbia S.C.: University of South Carolina Press.
- Tversky, A. 1977. Features of similarity. Psychological Review, Vol. 84, No. 4, 327-352. July 1977.
- Urban, W. 1939. Language and reality. London: G. Allen and Unwin, Ltd.
- Whorf, B. 1939. The relation of habitual thought and behavior to language. In J. Carroll (Ed.) Language, thought, and reality: Selected writings of Benjamin Lee Whorf, 134-159. New York: John Wiley & Sons. 1956.
- Winston, P. 1978. Learning by creating and justifying transfer frames. Artificial Intelligence, Vol. 10, 147-172.